Before the
FEDERAL COMMUNICATIONS COMMISSION
Washington, D.C. 20554

In the Matter of	)	
	)	
	)	
Developing a Unified Intercarrier Compensation Regime	)	CC Docket No. 01-92
	)	
	)	

DECLARATION OF LYMAN CHAPIN

I, Lyman Chapin, hereby declare and state as follows:

I. QUALIFICATIONS

1. I am the co-founder of Interisle Consulting Group, a firm that advises companies, non-profit organizations, and government agencies on matters of Internet technology and governance, and a Fellow of the Institute of Electrical and Electronics Engineers. For almost thirty years, I have played a leadership role in developing the network routing and interconnection architecture, protocols, and policy framework that support today's globally pervasive Internet. Prior to the creation of Interisle Consulting, I held a variety of technological positions including Chief Scientist at BBN (formerly Bolt, Beranek & Newman), which was one of the leading companies in the creation of the Internet and one of the first commercial Internet Service Providers. My responsibilities at BBN included the architecture and design of research, military, and commercial networks, and developing BBN's research results into business opportunities.

2. In addition, I was a principal architect of the Open Systems Interconnection reference model and protocols, which are both the generally accepted reference points for discussions of network architecture and the source of many of the technical innovations that have

been incorporated into today's Internet**.** I have served on many United States and international boards, associations, and committees responsible for establishing network and transport architecture, service, and protocol standards for the global Internet. For example, I was a member of the Internet Architecture Board from 1989 to 1993, and served as its Chairman for the last two years of that period. Between 1983 and 1999, I served as Chairman and in a variety of other positions for the Special Interest Group on Data Communication of the Association for Computing Machinery. I also served as a Director of the Internet Corporation for Assigned Names and Numbers; standards area director for the Internet Engineering Steering Group; co-founder and trustee of the Internet Society; United States representative to the networking panel of the North American Treaty Organization Science Committee; and United States representative to the computer communications technical committee of the International Federation for Information Processing.

3. Over the past 20 years, I have authored and co-authored a number of publications discussing issues related to Internet standards and protocols, including interconnection and the exchange of traffic. My publications include *Open Systems Networking: TCP/IP and OSI* (1993); *Routing Issues in Interconnecting IP Networks* (2000); and *Communication Systems: The State of the Art* (2002). A copy of my curriculum vitae and list of selected publications is attached as Exhibit A.

II. <u>SUMMARY</u>

4. I have been asked by Verizon to analyze the self-organized and self-regulating structures that govern interconnection and traffic exchange on the Internet today, and how those structures have evolved, and will continue to evolve, in the face of technological and market change. This discussion is based on my first hand experience with and knowledge of the

evolution of the Internet and the self-organized structures that govern it, as well as extensive reading in the literature of the field.

5. Almost every aspect of Internet technical development, deployment, and operation is determined by self-organizing, self-regulating groups that have proven, time and again, their ability to create and maintain technical, architectural, business, and governance policies and practices that encourage high-quality engineering, broad interoperability, and continued creation of value, while truly representing global consensus and thereby keeping participants on board. This approach—often loosely referred to as "the Internet approach"—is regularly cited as the key to the Internet's phenomenal success, because it enables the Internet to adapt quickly and efficiently to the rapid pace of change and innovation in telecommunications technology, operations, and public policy.

6. An important application of the Internet approach allows the many different corporate, institutional, and government entities that actually own and operate the Internet's mesh of individual networks to negotiate interconnection agreements with each other independently, without regulatory intervention. Internet Service Providers ("ISPs")[1] connect their separate networks to each other in order to exchange traffic between their customers and the customers of other ISPs, in accordance with a broad range of highly individual "peering" (paid and unpaid) and "transit" agreements[2] that specify the business and financial terms of their interconnection, including where and how traffic will be exchanged and whether and how compensation will be paid.

[1] In this declaration, I use the term "Internet Service Provider" to refer generally to any entity that provides Internet services.

[2] These terms are described in detail in Section IV.

7. Seamless Internet connectivity from any source to any destination is therefore the net result of countless independent decisions by individual ISPs concerning whether and how to interconnect. This approach capitalizes on the strong business incentives for ISPs to interconnect; no single ISP's network can reach every corner of the globe, and the market has shown that an ISP's interconnections with others is an important source of business opportunity. As a result, the Internet as a whole is always fully interconnected—the customers of every ISP can communicate with the customers of every other ISP, whether or not any particular pair of ISPs is connected. As studies by a wide variety of public and private organizations have repeatedly concluded, this market-driven model represents the most effective and efficient way to provide ubiquitous public Internet connectivity without being either anti-competitive or inequitable.

8. Moreover, the most effective and efficient way to ensure that the Internet continues to prosper in the future is to allow this market-driven interconnection model to continue. Top-down attempts to regulate, either in the name of "improving" the Internet itself, redressing perceived inequalities in access or pricing, or furthering policy objectives—no matter how well intentioned or carefully crafted—are contrary to the fundamental self-organizing, decentralized nature of the Internet, which is the most important source of the Internet's vitality. Such regulatory efforts necessarily run the risk of being destabilizing and harmful, and they should be rejected.

9. The remainder of this declaration provides additional information supporting these conclusions. Section III provides a brief history of how today's apparently seamless global Internet evolved from a few independent, unconnected networks. Section IV discusses interconnection in today's Internet, including a discussion of Internet Service Providers'

negotiated agreements to exchange traffic and the self-organized, non-governmental groups and associations that oversee network interconnection. In Section V, I conclude that the environment of non-regulation in which the Internet has developed has been essential to its ability to respond rapidly to technological change and shifting market demands with innovative and valuable new capabilities and services. In Section VI, I conclude that the most effective and efficient way to preserve the uniquely valuable properties of the Internet and to ensure that the Internet continues to prosper in the future is to defer to the "Internet Approach"—the self-governing, market-based environment in which the Internet now flourishes—and not to subject it to governmental regulation.

III. <u>The Origins of Internetworking, Interconnection, and the Internet</u>

10. The Internet as we know it today is the product of developments and innovations from a variety of sources that have evolved together over the past thirty years. Understanding the way in which the Internet evolved from earlier networks, and incorporated concepts of internetworking and interconnection over time, is essential to understanding the way in which the Internet works—not only from the perspective of technical architecture, but also from the perspective of operations, economics, and governance.

11. In the 1950s and early 1960s, long before Local Area Networks and Personal Computers, "computer communication" meant connecting input/output and storage devices (such as card readers, terminals, magnetic tape drives, and printers) to mainframe computers, typically within the same room. Early efforts to connect computers to each other led to "networks" based on a variety of different proprietary communications technology and protocols When there were just a few of these homogeneous networks, it was possible to exchange information between them relatively simply by building a device that could translate between one network's internal

protocols and another's. However, as the number of networks grew, the combinatorial complexity of connecting all of the networks to each other led to the idea of "internetworking." Internetworking is a technical architecture that enables networks based on different telecommunication technologies and protocols to exchange data, creating a "network of networks." Today, internetworking is the common operating mode throughout the Internet, which relies on the standard Internet Protocol, commonly referred to as "IP."

12. "Internetworking" is about technology. "Interconnection" refers to the operational and financial agreements that enable the owners and operators of different networks to collaborate as business entities in the use of internetworking to provide seamless end-to-end Internet connectivity to all of their individual customers. Today, interconnection takes place at hundreds of public and private exchange points at which two or more Internet Service Providers make technical, administrative, and economic arrangements to exchange traffic.

13. Remarkably, some of the earliest thinking about internetworking in the early to mid-1960s embraced the three key concepts that underlie the architecture of today's global Internet. These key concepts are:

- *Packet Switching*. Packet switching involves a distributed network of computers, each capable of exchanging and forwarding data, with (a) redundant links; (b) no central control; (c) messages broken into equal-sized "packets" at the source and re-assembled at the destination; (d) variable routing of packets depending on the availability of links and nodes, such that a series of packets from a given source to a given destination do not necessarily all follow the same path; and (e) automatic reconfiguration of routing tables after the loss of a link or mode.

- *Best-Effort Service*. Each computer in a packet-switching network attempts to forward any packet it receives that is addressed to a computer other than itself. Most of the time, forwarding will be successful. Occasionally, due either to technical flaws or to contention for the same link by multiple computers, it will fail. But the failure of a packet to reach its destination does not interrupt the overall flow of messages: the receiving computer simply requests re-transmission of the missing packet and waits to receive it prior to re-assembling and delivering the message
- *Application Independence*. The network should be adaptable to any purpose, whether foreseen or unforeseen, rather than engineered specifically for a single purpose.[3]

A. Early Government Networks

14. Most of the early work on computer communication in the 1960s was sponsored by the United States Department of Defense Advanced Research Projects Agency[4], which funded 17 incompatible time-sharing and interactive computing[5] projects before deciding, in 1968, to commission the design of a distributed communications network that would connect the many different Advanced Research Projects Agency sites around the country. This network, known as the “ARPAnet,” began operating in 1969, initially connecting just four university research laboratories.

---

[3] The public switched telephone network is an example of a network that was originally purpose-built for a particular application: analog voice communication.

[4] Today this agency is called the Defense Advanced Research Projects Agency.

[5] In a time-sharing computer system, more than one person can use the computer at the same time. In an interactive computer system, a user can interact with the computer in “real time” through a terminal interface.

15. At the outset, in 1969, the ARPAnet was not an "internet" as we think of it today—each of its four computer hosts was connected to an Interface Message Processor by a proprietary serial link and protocol, and the Interface Message Processors communicated with each other over 56 kilobit per second lines leased from the telephone company, using an ARPAnet-specific "host-to-host protocol" that was referred to as the Network Control Program. Other packet networks, based on other protocols, were being developed at the same time. The first papers describing "packet network interconnection" were published in 1973; the ARPAnet began using the standard Internet Protocol in 1977. The Internet Protocol became a Military Standard—mandated for all Department of Defense-funded networks, including the ARPAnet—in 1983.

16. From its inception, the ARPAnet was managed by an informal and mostly self-selected group of engineers and managers who designed, installed, and operated the network, known as the Network Working Group. The tradition of self-management by the people designing, installing, and operating the network was established at the very first Network Working Group meeting, and has carried through to the structures that oversee the Internet today—particularly the Internet Engineering Task Force and the various Network Operators Groups discussed in more detail below.

17. The ARPAnet proved to be immensely useful to the defense community. Connections to the ARPAnet, however, were controlled by an Acceptable Use Policy, which restricted the use of the ARPAnet to a closed community of Department of Defense employees and defense contractors, and limited the ways in which it could be used, such as to conduct or support research funded by the Department of Defense. Other agencies therefore developed their own similar networks. For example, the United States Department of Energy created MFENet

for its magnetic fusion energy researchers and HEPNet for its high energy physicists, and the National Aeronautics and Space Administration ("NASA") established the SPAN network for its space physicists. Each of these agency networks was governed by its own Acceptable Use Policy, limiting its use to the agency's own community of funded researchers.

B. Early Non-Government Networks

18. The restrictive Acceptable Use Policies that governed the early government networks led to the development in the 1970s and early 1980s of a number of other, non-governmental, networks, as computer scientists who could not connect to one of the government-agency sponsored networks established alternative networks for their own use. Two such networks that feature prominently in the Internet's history were USENET and BITNET. Although USENET and BITNET were not limited by the same types of Acceptable Use Policy restrictions as the government-sponsored networks, as a practical matter they were established for, and used by, the academic and industrial computer science community (rather than the public at large).

19. Like those operated by government agencies, these networks were not regulated by an outside authority, but rather self-regulated by the engineers and managers who operated them.

C. Beginnings Of Interconnection

20. The first step toward the interconnections that would lead to today's Internet was a response to the practical awkwardness of operating multiple non-communicating networks, which led to the establishment of two federally-funded exchange points for the networks operated by NASA, the Department of Energy, and the Department of Defense: the Federal Internet Exchanges at the University of Maryland and at NASA's Ames Research Center in

Mountain View, California. Following the tradition that was by this point well-established, these interconnection points were managed by two informal groups of engineers and managers: the Federal Networking Council, which handled administrative matters, and the Federal Engineering Planning Group, which handled technical matters.

21. Interconnection took a second step forward with CSNET, a project created in 1981 under a grant from the National Science Foundation to link all of the computer science departments and industry labs engaged in computing research. CSNET provided Transmission Control Protocol/Internet Protocol interfaces with USENET, BITNET, and the X.25[6] networks, and established nameserver databases to enable any computing researcher to locate any other.

22. The development of CSNET highlighted the disconnect between the "haves" and the "have nots" in the computing research community—between those who could find a government agency or contractor to sponsor their connection to the ARPAnet, and those who could not (connecting instead to CSNET). In modern terms, we would say that the customers of one ISP (ARPAnet) could not communicate with the customers of another ISP (CSNET), because no mechanism existed to reconcile the different Acceptable Use Policies of the two networks.

23. This disconnect persisted until the CSNET managers came up with the idea that we now call "peering," or interconnection without explicit accounting or settlement.[7] A landmark agreement between the National Science Foundation and the Department of Defense Advanced

[6] The Consultative Committee on International Telephone and Telegraph, a United Nations treaty organization that is today called the International Telecommunications Union, created a number of standards for packet-switching networks in the late 1970s and early 1980s, among them a standard referred to as "Recommendation X.25."

[7] The concept of peering is described in greater detail in Section IV. As the term is used here, it refers to "unpaid peering"; the concept of "paid peering" did not exist before the transition to private (commercial) ownership and operation of Internet networks which is described in the next section.

Research Projects Agency allowed National Science Foundation grantees and affiliated industry research labs access to the ARPAnet, as long as no commercial traffic flowed through the ARPAnet. This agreement was the turning point at which the evolution of commercial network interconnection began.

24. In 1985, the National Science Foundation funded the first five national supercomputer centers, and its plan included a high-speed national backbone network—the NSFnet—to connect them. Over the next few years, the National Science Foundation commissioned the development of a deliberate architecture of backbones and regional networks that introduced the idea of hierarchy into the Internet topology. By 1990, the NSFnet had become the backbone of the modern Internet.

D. Privatization, Commercialization, and Globalization of the Internet

25. In 1993, as the National Science Foundation began the transition to private ownership and management of the NSFnet, it established four geographically distributed, privately owned and operated Network Access Points, operated by Sprint, Pacific Bell, Ameritech, and Metropolitan Fiber Services. Under the terms established by the National Science Foundation, a Network Access Point operator was required to provide and operate an interconnection facility on a nondiscriminatory basis, using published pricing and established technical operating specifications. These Network Access Points were the first commercial Internet exchange points, where any interested party could co-locate equipment and connect its network to the NFSnet backbone or to other networks. As the original Network Access Points (also referred to as Metropolitan Area Exchanges) became increasingly congested, many network providers began creating their own private Network Access Points, which further extended the commercial Internet exchange model.

26. In 1996, the National Science Foundation handed over the management of its backbone to commercial ISPs. End users connected to ISPs by placing calls over the public telephone network to modem banks operated by the ISPs, or via leased circuits of higher capacity. ISPs, in turn, connected to regional or backbone networks at Network Access Points or exchange points.

27. This interconnection hierarchy did not, however, correspond to a strict hierarchy of ISPs and backbone providers as business entities. Some providers were vertically integrated, operating in every business from high-capacity backbone traffic down to dial-up lines. Others specialized in providing one form or another of connectivity to one or more specific markets.

28. The economic incentives and tradeoffs that are so richly diverse in today's Internet began to develop as soon as commercial ISPs recognized that their interconnection arrangements could be a source of competitive advantage. Any ISP could connect to one of the public Internet exchange points, but the opportunity to achieve better performance, particularly for destinations that would be several "hops" [8] away using a public exchange, led many ISPs to explore direct interconnection of their networks with those of other ISPs. The growing number of ISPs, and the variety of different ways in which the rapidly expanding Internet services market drove the development of creative combinations of public and private ISP interconnection, ensured that the Internet as a whole would always be fully interconnected; the customers of every ISP could communicate with the customers of every other ISP, whether or not any particular pair of ISPs installed an explicit public or private interconnection.

[8] In Internet routing parlance, a "hop" is the distance between one router and the next along the path that a data packet follows from source to destination. In general, the more ISPs and exchange points are involved in creating that path, the more "hops" a packet will travel to get from source to destination.

29. As the number and diversity of Network Access Points increased, the potential complexity of hundreds or thousands of ad-hoc bilateral arrangements pointed to the need for an overarching, neutral policy framework within which providers could implement mutually beneficial cost-sharing interconnection agreements. It was at this juncture that there began to emerge a large number of privately operated Network Access Points, also known as Internet exchange points, which provided a uniform set of technical and administrative services (*e.g.*, interconnection, traffic routing based on sophisticated criteria, operational support of routing equipment, traffic metering, billing, and clearing and settlement of charges between parties). These exchanges provided a framework that allowed multiple providers of different sizes, scopes, and operating philosophies, serving the same or different markets, to interconnect in ways appropriate to each.

30. The Internet developed earlier and more rapidly in North America than in the rest of the world. Initially, as discussed in more detail below, the greater economic and regulatory maturity of Internet interconnection in North America produced remarkably inefficient configurations in which non-North American ISPs—even those in the same or adjacent countries—connected to each other through exchange points in the United States. By the mid-1990s, however, consumer and content-provider interest in the Internet outside of North America had generated more than enough market incentive to redress this imbalance, and a robust global market in interconnection and exchanges emerged.

IV. Self-Management and Interconnection in Today's Internet

31. ISP interconnection involves a number of operational, administrative, financial, and legal issues that go beyond the simple exchange of network traffic, all of which require cooperation and collaboration among multiple ISPs. Examples include:

- securely exchanging interdomain routing information;
- providing services that work across multiple ISPs;
- detecting and responding to denial of service attacks (and possibly other forms of distributed, multi-ISP attack that have yet to be seen);
- controlling spam, phishing, and other intrinsically multi-ISP exploits; and
- enforcing national public policy mandates, such as universal service, emergency warnings, wiretapping, and the like.

32. Interconnecting all of the ISP networks into the global Internet depends on technical and procedural standards; in a narrower sense, the interconnection of any two particular networks depends on establishing the operational and financial terms (including compensation) for that particular interconnection. The tradition of self-management by the designers, installers, and operators of the network that first began in the late 1960s with the ARPAnet and the Network Working Group has carried through to these decisions today. All aspects of the operation of the global Internet—from establishing industry-wide standards to setting the terms and prices for the interconnection of two particular networks—are decided by the people who operate the networks that make up the Internet itself.

A. Self-Management of Operational Aspects by Industry Groups

33. The standard-setting and decision-making necessary to enable networks to join together in the Internet are accomplished through self-organized industry groups. For example, technical standards are developed by a self-organized group of Internet engineers and computer scientists known as the Internet Engineering Task Force:

> "…a loosely self-organized group of people who contribute to the engineering and evolution of Internet technologies. It is the principal body engaged in the development of new Internet standard specifications. The IETF is unusual in that

> it exists as a collection of happenings, but is not a corporation and has no board of directors, no members, and no dues."[9]

The Internet Engineering Task Force is closely associated with another self-organized group, the Internet Architecture Board, and is housed, administratively, within yet another such group, the Internet Society. The "loosely self-organized" Internet Engineering Task Force and related organizations have proven, over a 20 year history, to be effective at establishing workable standards and highly adaptive to the rapid growth and change that have occurred within the Internet.

34. Self-governance also dominates the establishment of operational standards and practices for ISPs. Since the earliest days of the Internet, the operators of interconnected networks have met both informally and formally to share technical information and coordinate operating principles and practices. In 1994, members of the former NSFnet "regional-techs" meetings[10] formed an expanded group, called the North American Network Operators Group, with the mission of promoting and coordinating the interconnection of networks within North America and to other continents, serving as an operational forum for the coordination and dissemination of technical information related to backbone and enterprise networking technologies and operational practices. The North American Network Operators Group has been highly effective in enabling ISPs and backbone providers to coordinate their activities to efficiently provide seamless service to a broad market. Other regions of the world have followed the Operators Group model and have developed or are developing similar groups, such as the

---

[9] *Tao of the IETF—A Novice's Guide to the Internet Engineering Task Force*. Susan Harris and Paul Hoffman, October 2004. <http://edu.ietf.org/tao>.

[10] The "regional-techs" meetings, which began in 1985, were informal periodic meetings of the technical community of engineers and managers that operated the backbone and regional networks of the NSFnet.

African Network Operators Group, the South Asian Network Operators Group, and the Pacific Network Operators Group.

35. In the Internet, self-organized groups also perform the important governance functions of promoting an efficient exchange of value in the allocation of resources. The Internet depends on two main types of resource:

- Physical, tangible infrastructure such as communications links and switching facilities, which are owned by the private parties that own networks; and
- Virtual resources, such as Domain Names and IP Addresses, which are allocated by industry-based organizations.

Domain names, such as "coca-cola.com" or "lightbulbs.com," combine aspects of traditional intellectual property (*i.e.*, trademarks and service marks) with the technical infrastructure required to cause the names to perform their intended function. Numeric *IP Addresses* uniquely identify each computer connected to the Internet. The Internet Engineering Task Force defines these naming and addressing schemes, which are then administered by the Internet Corporation for Assigned Names and Numbers. The Internet Corporation for Assigned Names and Numbers oversees the allocation of Domain Names, which is implemented by a globally distributed hierarchy of domain name registry and registrar companies; the allocation of IP Addresses, which is implemented by five Regional Internet Registries; and the operation of the decentralized mechanism (the Root Server System) whereby Domain Names are translated into IP Addresses.

B. Interconnection

36. The same self-organizing approach that governs these broad, Internet-wide decisions also extends to the decisions made by individual ISPs to interconnect (or not) and the terms under which they will do so. ISPs negotiate agreements that specify the terms of their interconnection, including where and how traffic will be exchanged and whether and how

compensation will be paid. Different interconnection agreements may contain vastly different terms, depending on the needs and value assessments of the particular ISPs involved.

37. At its most basic, an interconnection agreement says "You carry some traffic for me, in return for which I'll do something—either carry traffic for you, or pay you, or some combination of the two." Indeed, from a purely technical standpoint, ISP interconnection is no more complicated (or controversial) than simple internetworking, in which routers connected by communication links of various kinds compute routes through the Internet based on information they have received from hosts (end users) on any networks to which they are directly connected and from other routers. One way in which two networks may interconnect is at a third-party Internet exchange point. In its simplest form, such an Internet exchange point is a physical place (typically a room in a building) in which Internet routers are installed. ISPs that want to use the exchange point to connect to other ISPs run one or more links from their own routers to the exchange point, where they connect those links to the exchange point routers. The ISP routers and the exchange point routers exchange information about where different groups of Internet hosts—identified by their IP addresses—are located, using routing protocols such as the Border Gateway Protocol. ISP A might learn, for example, that a group of Internet users who are customers of ISP B can be reached through an exchange point to which both A and B are connected, and decide to use the exchange point to reach those users. Traffic from users on A's network to users on B's network would flow over A's network as far as the exchange point, where it would pass to B's network.[11] Alternatively, two ISPs may decide to connect their networks directly to each other, rather than meeting at a third-party exchange point.

[11] In practice, of course, the way in which traffic flows are managed at exchange points is more complicated than in this example. Nevertheless, this example captures the essence and basic underlying process of such exchanges.

38. In reality, however, interconnection agreements are not so simple. They are often tailored very carefully and minutely to the specific circumstances of the parties involved, particularly when those parties are large ISPs. The term "interconnection policy" refers to the way in which the technical and contractual arrangements that ISPs negotiate with each other to interconnect directly or at Internet exchange points are influenced by the business objectives and policies of each of the parties. The term "interconnection economics," on the other hand, refers to the way in which interconnecting ISPs assess and manipulate the economic variables that determine the viability of interconnection as a business proposition.

39. The interconnection arrangements made by the heterogeneous mix of ISPs that operate today's Internet—large and small; local, national, and global; public and private—can be grouped into four basic models:

- *Bilateral settlements.* Two ISPs interconnect. Each accepts traffic from the other that is destined for its own customers. Neither delivers traffic to third parties on behalf of the other. Each charges for the volume of traffic it accepts from the other.

- *Sender Keep All.* As with bilateral settlement, two ISPs interconnect, but the two network operators agree to exchange traffic without explicitly charging or otherwise accounting for it.[12]

- *Transit.* One ISP, the transit provider, accepts traffic originating within the other ISP's network, destined not only for the provider's own customers but also for third party networks with whom the provider in turn connects. The provider charges a fee for carrying the other ISP's traffic.

- *Multilateral exchanges.* An ISP connects to an Internet exchange point, at which traffic is routed to other ISPs' networks via equipment provided by the exchange and according to rules administered by the exchange. The exchange calculates a

[12] This arrangement is called "sender keep all" because each party bills its own customers for the traffic they send and keeps all of the resulting revenue (rather than paying some of it to the other party as compensation for carrying the traffic over its network); it is also referred to as "bill and keep."

net settlement due to or from each ISP based on the volume of traffic that each ISP carries on behalf of others.[13]

The "sender keep all" model is often referred to as "peering"; the "bilateral settlements" model is often referred to as "paid peering."[14]

40. In each case, the interconnection agreement is based on a *perceived equitable exchange of value* between the two interconnecting parties. If both parties believe that the fact of interconnection itself results in an equitable exchange of value, they may agree to peer without cash payment or other compensation. On the other hand, if either party believes that the benefits of interconnection are unequal (for any of the reasons described below), the ISP that stands to gain more from the arrangement will generally be required to compensate the other. ISPs base their peering decisions—whether or not to interconnect with another ISP, and whether or not cash or non-cash compensation should be required (and if so, how much)—on many factors,[15] including:

- *Geographic coverage*. ISPs may prefer partners with networks that provide overlapping geographic coverage, in which case an interconnection agreement would be symmetrical; or they may prefer partners with networks that cover different geographical areas, in which case interconnection would extend each network's geographic reach.

[13] Many exchanges operate in the current competitive market; ISPs can choose which (if any) exchanges to use, and choose those that offer the most advantageous combination of technical, economic, and contractual features.

[14] The term "peering" is occasionally used more loosely to refer to either of the models in which two ISPs agree to carry traffic to and from each other's customers, whether or not compensation is part of the agreement. Used in this way, the term "peering" distinguishes bilateral agreements between ISPs at roughly the same level in the interconnection hierarchy from fee-for-service agreements between a customer ISP at a lower level in the hierarchy purchasing a "transit" service from a provider ISP at a higher level.

[15] Although ISPs were once largely unwilling to reveal their interconnection policies, today competition forces them to publish their policies openly. The fact that ISP interconnection policies have become increasingly public speaks to an increasingly transparent and participatory market.

- *Technology*. ISPs may preferentially choose or reject interconnection partners based on support for technical standards or access to a desired technology.

- *Operations*. ISPs may require a certain level of operational support from interconnection partners.

- *Routing*. ISPs may require that interconnection partners support specific routing policies and practices.

- *Size*. ISPs may decide to connect their networks directly only to networks of a similar size, and to require (or pay) compensation to interconnect with networks that are smaller (or larger).

- *Anticipated traffic symmetry*. An ISP may choose to peer on a sender-keep-all basis only with networks that deliver to the ISP roughly[16] the same amount of traffic that they accept from it, and to require (or pay) compensation to interconnect when traffic volume to and from the other network is expected to be highly asymmetric.

- *"Blacklisted" behavior*. Some ISPs will refuse to carry traffic that originated with another ISP that has been "blacklisted" for sponsoring spam or phishing attacks (or other misbehavior).

Additional, highly individual factors not apparent to an outside observer may also apply to an ISP's calculation of the value of a particular interconnection. For example, if one ISP's network geography, customer mix, or traffic mix reinforces an important element of another ISP's business strategy, the other ISP's perceived value of interconnection might be higher.[17]

V. <u>The Internet Has Flourished under—and because of—the Market-Based and Non-Regulatory "Internet Approach"</u>

41. Almost every aspect of Internet technical development, deployment, and operation is governed by a self-organized, self-regulating structure. This approach—often

[16] The anticipated symmetry of traffic volume between two ISPs is generally calculated as a range averaged over some period of time; *e.g.*, "+/– 5% over 24 hours." Many sender-keep-all peering agreements are designed to convert to paid peering if either party exceeds the specified threshold.

[17] Because so many factors affect each interconnection decision, it is extremely difficult to analyze the economics of any particular interconnection arrangement using external, objective criteria.

loosely referred to as "the Internet approach"—is regularly cited as the key to the Internet's phenomenal success. What has given vitality to the Internet approach is not simply that it meets the needs of the current environment; it is the fact that it relies on underlying *processes* that are flexible, adaptive, and efficient. The Internet approach is driven by self-organizing, self-regulating groups that have proven, time and again, their ability to create and maintain technical, architectural, business, and governance policies and practices that encourage high-quality engineering, broad interoperability, and continued creation of value, while truly representing global consensus and thereby keeping participants on board. As a result, self-regulation has allowed the Internet to adapt quickly and efficiently to the rapid pace of change and innovation in telecommunications technology, operations, and public policy.

42. The Internet approach has succeeded, in part, by capitalizing on networks' strong incentive to interconnect. These incentives began to develop as soon as commercial ISPs recognized that their interconnection arrangements could be a source of competitive advantage. Any ISP could connect to one of the public Internet exchange points, but the opportunity to achieve better performance, particularly for destinations that would be several "hops" away using a public exchange, led many ISPs to explore direct connection of their networks to those of other ISPs. It also led to the proliferation of Internet exchange points, as exchange point operators competed to provide the most advantageous environment for multiple ISPs to interconnect.

43. Architecturally, the Internet is "self-healing"—its routing protocols explore multiple ways to get from one point to another, using any available path through as many different links and routers as necessary, automatically bypassing traffic bottlenecks, outages of equipment or circuits, and other obstacles. This dynamic route-finding property also allows the

Internet to support communication between any two Internet users, whether or not they are connected to the same ISP, and whether or not their ISPs are directly connected to each other. For example, if an Internet user in Minneapolis, connected to the local ISP "TwinCities.net," wants to visit the web site of a bank in New York that is connected to Verizon, her data might flow through an exchange point at which TwinCities.net and Verizon connect directly to each other. However, as TwinCities.net is (presumably) much smaller than Verizon and its network extends over a much smaller geographical area, the two ISPs may not directly connect to each other anywhere. In this case, the user's data might travel first to a connection between TwinCities.net and "Midwest ISP" in Minneapolis; over the network owned by Midwest ISP to St. Louis, where Midwest ISP connects to Verizon; and then over Verizon's network to the bank in New York. This routing happens transparently, dynamically, and in real time. Neither the user in Minneapolis nor the bank in New York knows or cares that it is happening.

44. The topology of Internet interconnection has emerged over the past decade as an important factor in studies of Internet resilience and survivability. A corollary to many of these studies is the observation that the self-healing properties of the Internet architecture guarantee that the Internet as a whole would remain fully interconnected even if most of the direct connections between individual ISPs were removed. The fear of Internet "balkanization" as a result of large ISPs refusing to interconnect with smaller ISPs is, in today's Internet, completely unfounded. One ISP, or even a group of ISPs operating as a cartel, would find it virtually impossible to isolate another ISP or hold up its traffic. As traffic will in any case be routed around such an "obstruction," an economically self-interested ISP has an intrinsic motivation to connect, so as to participate in the revenue-generating potential of that traffic.

45. In spite of industry consolidation, the backbone transit business is becoming more competitive rather than less so, as suggested by comments such as the following:

> "Trends in transport pricing over the past six months have created a disruptive change by lowering the barriers for small and regional networks to develop robust national backbones for application delivery, peering, network performance and business expansion."[18]

46. A specific example of the ability of market forces to promote efficient outcomes for Internet connectivity is found in the experience of non-North American ISPs. Because the Internet developed earlier, and more rapidly, in North America than in other parts of the world, the interconnection arrangements between North American ISPs and networks in other countries initially were biased strongly in favor of the North American ISPs. Until relatively recently, it was common for Internet users in Taiwan, for example, to communicate with other Internet users in Japan or Singapore over a path that led through an exchange point in California (MAE-West), with the Asian network operators paying the full cost of the trans-Pacific links. This imbalance arose both from the early absence of an exchange infrastructure in other parts of the world and from the much more favorable (largely unregulated) economics of Internet telecommunications in North America, which meant that even where a link existed between, for example, Germany and France, the cost of connecting through an exchange point on the east coast of the United States—even including the cost of trans-Atlantic transport—could be substantially lower than the cost of a direct connection.

47. As recently as five years ago, ISPs in non-North American countries who were unhappy with this imbalance sought to correct it by attempting to force North American ISPs to subsidize the cost of inter-regional links. However, as dozens of viable regional Internet

[18] Jay Adelson, founder and CTO, Equinix; conference session announcement, 2005 ISPCON.

exchanges emerged outside of North America,[19] the pressure to regulate international ISP interconnection in favor of non-North American ISPs substantially evaporated. Market forces now drive ISP interconnection decisions in many other countries as effectively as they do in North America.

VI. Continued Self Management, Rather than Governmental Regulation, Will Ensure the Internet's Continued Evolution

48. Today's Internet is the way it is because of the way it developed. In every arena—technical standards, operating practices, and resource allocation—policy is established by self-organized, inclusive organizations, operating with a high degree of transparency, and representing a broad constituency.

49. This approach is nearly inevitable, given the inherently decentralized native architecture of the Internet and the heterogeneous, global market in which it operates. The incentives are well aligned: due to the network effect, continued growth of the Internet is a rising tide that lifts all boats, which creates a strong bias toward policies that facilitate growth and efficiency.

50. The growing number of ISPs and Internet exchange points, and the variety of different ways in which the rapidly expanding Internet services market has driven the development of creative combinations of public and private ISP interconnection, ensures that the Internet as a whole is always fully interconnected; the customers of every ISP can communicate with the customers of every other ISP, whether or not any particular pair of ISPs has installed an explicit public or private interconnection. Because the Internet is architecturally self-healing, the

[19] A current list of Internet exchange points is maintained at https://www.peeringdb.com/private/exchange_list.php; at the time at which this declaration is made, 67 of the 99 listed exchanges are located outside of North America.

possibility that an ISP could find itself unable to connect its customers to some part of the Internet because one ISP (or even many ISPs) refused to interconnect with it is vanishingly small; there are simply too many available connection points, public and private, and the architecture of the Internet ensures that traffic will flow end-to-end regardless of where an ISP is connected. The architecture of the Internet and the competitive diversity of the ISP market jointly ensure both low barriers to entry for ISPs of all sizes and an Internet marketplace that cannot be controlled by even the largest ISPs.

51. The Internet will face many new challenges in the future. New application services may change the nature of interconnection and the content of interconnection agreements, as ISPs seek to differentiate themselves by offering new services, or guarantees of specific network performance, only to their own customers. This is unlikely, however, to result in "balkanization" of the Internet or a loss of full and complete connectivity, because there is no reason to expect that the market forces now driving a high degree of interconnection, as discussed above, will change significantly. Furthermore, responding to new technology is well within the scope of the "Internet approach": the existing policy mechanisms are well equipped to adapt to these changes, as they have to equally disruptive challenges and changes in the past.

52. By contrast, external regulation, no matter how well intentioned or carefully crafted, is contrary to the fundamental, decentralized nature of the Internet—the most important source of its vitality—and runs the risk of being destabilizing and harmful. Today, the self-organized, self-regulating Internet is thriving. In the absence of such a dramatic and obvious change, regulators should allow the "Internet approach" to operate without interference.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Dated: _5/23/05__

A. Lyman Chapin
Lyman Chapin

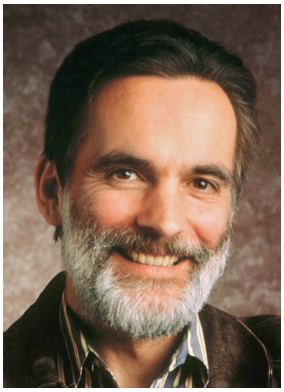

LYMAN CHAPIN
Partner
Interisle Consulting Group
+1 617 686 2527
lyman@interisle.net

Profile

Lyman Chapin is a co-founder of and partner at Interisle Consulting Group, LLC, where he advises companies, non-profit organizations, and government agencies on telecommunications network resilience and critical infrastructure protection, and serves as an expert advisor on matters of Internet technology and governance. He was most recently Chief Scientist at NextHop Technologies; before joining NextHop in 2001, he was Chief Scientist at BBN Technologies. Mr. Chapin is a Fellow of the IEEE, and was a founding trustee of the Internet Society. He has served as a Director of the Internet Corporation for Assigned Names and Numbers (ICANN) and as chairman of the Internet Architecture Board (IAB), the ACM Special Interest Group on Data Communication (SIGCOMM), and the ANSI and ISO standards groups responsible for Network and Transport layer networking standards. Mr. Chapin was a principal architect of the Open Systems Interconnection (OSI) reference model and protocols, and is the co-author of *Open Systems Networking—TCP/IP and OSI*. His professional interests include Internet technology, particularly internetwork routing, traffic engineering, and the Domain Name System (DNS); Internet governance; information security and personal privacy; and identity management. He lives with his wife and two daughters in Hopkinton, MA (USA).

Education

B.A. Mathematics, Cornell University, 1973

Professional Experience

Chief Scientist, NextHop Technologies (Apr 01 – Jul 02)
Co-founder and technical visionary for the strategic evolution of NextHop's Internet infrastructure business at the intersection of routing, traffic engineering, and network management and control.

Chief Scientist, BBN Technologies (Jan 98 – Mar 01)
Responsible for strategic technology initiatives that combined the strengths of BBN with those of the other units of the Verizon Technology Organization (Verizon Laboratories and Federal Network Systems), particularly in the areas of Internet architecture and technology, wireless mobile networking, information security, and electronic payment systems.

Chief Technology Officer, BBN Professional Services (Apr 93 – Jun 99)
Responsible for creating and promoting a high level of technical sophistication and expertise as the distinguishing features of Professional Services engagements with financial industry and other commercial clients, and for the coordination of GTE CyberTrust (security) and Professional Services business and technical strategies.

Chief Network Architect, BBN Communications (Sep 90 - Mar 93)
Chief architect for BBN Communications hardware and software products. Responsible for technology development and marketing relationships with other BBN divisions and other companies, and for the management of BBN's participation in industry standards activities.

LYMAN CHAPIN
Partner
Interisle Consulting Group
+1 617 686 2527
lyman@interisle.net

Professional Experience (continued)

Senior Consulting Engineer, Data General Corporation (Oct 77 - Sep 90)
Founding member of DG's original Networks development group, responsible for network system design and the development of software products for distributed resource management (operating–system embedded RPC), distributed database access, X.25–based local- and wide-area networks, and OSI–based transport, internetwork, and routing functions in local- and wide-area networks. Director of the distributed systems architecture group, responsible for the development of DG's Distributed Application Architecture (DAA) and its directory and management services.

Systems Analyst, Systems & Programs Ltd., New Zealand (Aug 74 - Sep 76)
Programmer and systems analyst for large–scale mainframe–based business software development projects.

Professional Affiliations

IEEE
Fellow (Jan 00 – Present)

Internet Architecture Board (IAB)
Chairman (Jul 91 – Mar 93)
Member (Jan 89 – Mar 93)

Internet Society
Founding member of the Board of Trustees (Jun 92 – Sep 94)
Internet Society 1998 conference program committee co-chair

Internet Engineering Steering Group (IESG)
Internet Standards Area Director (Mar 93 – Jun 94)

Internet Corporation for Assigned Names and Numbers (ICANN)
Director (Oct 01 – May 04)

ISO/IEC JTC 1/SC 6 (Telecommunications standards)
Head of USA national delegation (Apr 84 – Nov 93)
Editor for ISO/IEC 8348 (Network service definition), 8473 (CLNP), 9542 (ES-IS routing protocol), and related Network and Transport standards

Artech House Publishers
Editor of the Communications Engineering series (2003 – Present)

Wiley Computer Publishing
Founder and editor of the Wiley Networking Council series (1997 – 2002)

NATO Science Committee
USA representative to the Networking Advisory Panel (Jan 00 – Jun 04)

ANSI task group X3S3.3 (Network and Transport layer standards)
Chairman (Feb 82 – Nov 93)

IEEE/ACM
Co-founder of the refereed journal *IEEE/ACM Transactions on Networking* (1993)
Inter-society steering committee for *IEEE/ACM Transactions on Networking*
(Jun 93 – Jun 97; chairman Jul 93 – May 95)

LYMAN CHAPIN
Partner
Interisle Consulting Group
+1 617 686 2527
lyman@interisle.net

Professional Affiliations (continued)

CommerceNet
Director (Jun 99 – Nov 03)

ACM Special Interest Group on Data Communication (SIGCOMM)
Chairman (Jul 91 – Jun 95)
Vice-chairman (Jul 85 – Jun 91)
SIGCOMM Conference program committee (1990-5, 2001)
SIGCOMM Conference general chairman (1999)
Mobicom 1995 conference steering committee
Latin America Networking Conference steering and program committees (2003)

International Federation for Information Processing (IFIP)
USA (ACM) representative to TC6 (Communication Systems) (Apr 94 – Present)
IFIP SmartNet 1999 conference program committee
IFIP Networking 2000 conference program committee
World Computer Congress 2002 program committee co-chairman

Selected Publications

Communication Systems: The State of the Art, IFIP 17th World Computer Congress. A. Lyman Chapin. Kluwer, 2002.

"Routing Issues in Interconnecting IP Networks." *Proceedings of the Networks 2000 Conference* (September 2000).

Integrated Network Management IV. Chapter 58, "Managing Complex Systems." A. Lyman Chapin. Chapman & Hall, 1995.

Open Systems Networking: TCP/IP and OSI. A. Lyman Chapin and David M. Piscitello. Addison–Wesley, 1993.

Internet System Handbook. Chapter 17, "The Billion-Node Internet." A. Lyman Chapin. Addison-Wesley, 1993.

Co-editor for two Feature Topic issues of *IEEE Communications* on "The Global Internet," May 1997 and June 1997.

"The State of the Internet: 1995." *Telecommunications* 29:1 (January 1995).

"The State of the Internet: 1994." *Telecommunications* 28:1 (January 1994).

"Introduction to Routing." *Connexions Magazine* 7:9 (September 1993).

"The Global Internet: The 'Academic Toy' Comes of Age." *Telecommunications* 27:5 (May 1993).

"The Internet Standards Process." RFC 1602 (March 1994).

"The Internet Architecture Board and the Future of the Internet." *EDUCOM Review* 27:5 (Sept/Oct 1992).

"Standards for Bridges and Gateways." *IEEE Network* 2:1 (January 1988).

"Internetworking in a LAN/WAN Environment." *Data Communications* (May 1986).

"An International Internetwork Protocol Standard." *Journal of Telecommunication Networks* (Fall 1984).

"Connections and Connectionless Data Transmission." *Proceedings of the IEEE* 71:12 (December 1983).